


1663

Computing in the Fast Lane
Magnetic Fields of Dreams
MCNP: A Code in Demand
Agent-Based Models



About Our Name: During World War II, all that the outside world knew of Los Alamos and its top-secret laboratory was the mailing address—P. O. Box 1663, Santa Fe, New Mexico. That box number, still part of our address, symbolizes our historic role in the nation's service.

Located on the high mesas of northern New Mexico, Los Alamos National Laboratory was founded in 1943 to build the first atomic bomb. It remains a premier scientific laboratory dedicated to national security in its broadest sense. The Laboratory is operated by Los Alamos National Security, LLC, for the Department of Energy's National Nuclear Security Administration.

About the Cover: Ross McDonald of the National High Magnetic Field Laboratory-Los Alamos prepares the 100-tesla multi-shot magnet for an experiment. One of the most powerful, repeat-use magnets in the world, the 100-tesla is opening up new frontiers in the study of complex materials.



LOS ALAMOS ARCHIVE

The main guard gate to the technical area at Los Alamos during World War II.



From Terry Wallace

Information Science and Technology Revolution

Computer simulation of complex systems is a hallmark capability at Los Alamos National Laboratory.

It began with John von Neumann, who joined the Manhattan Project in 1943. His very-difficult calculations concerning nuclear explosions made him realize that machines could be used to solve the numerical formulations of partial differential equations. This ultimately led to the Laboratory's first supercomputer: the Mathematical Analyzer, Numerical Integrator and Computer, affectionately known as MANIAC.

Ever since then, Los Alamos has pushed the frontiers in computers, building and operating the world's fastest machines, including the Cray-1 in the 1970s, Thinking Machines' CM-5 in the 1980s, and later this year, in collaboration with IBM, the first petaflop computer (a million billion operations per second)—Roadrunner.

Roadrunner increases performance through a dramatically new strategy—mixing different types of processors that can be optimized for different types of calculations. This strategy holds the promise of computers hundreds of times more powerful than Roadrunner within the next decade.

Hardware is only a small part of simulating


complex systems. Stanislaw Ulam, another Manhattan Project member, had the idea of evaluating complex processes statistically by what is now called the Monte Carlo method. Legend has it that while playing solitaire, Ulam thought about calculating the probability of winning, but realized that playing many games and counting the number of successful plays would give the answer more easily. He then realized that the new fast computers could use that strategy for problems of neutron diffusion and other processes described by certain differential equations, provided those processes could be translated into an equivalent form composed of a succession of random operations.

Ulam's idea gave birth to one of the Laboratory's most-powerful computer simulation tools, the Monte Carlo code MCNP.

Today, Los Alamos has world-class capability in simulation science and a focus on information science and technology (IST), reflecting the growing need to solve data, information, and knowledge issues in 21st-century science. From satellite-based sensors to medical imaging devices and intelligence databases, automated collection of terascale data is becoming standard. The generation of petascale simulation data is close at hand.

This edition of 1663 is a window into the IST revolution to come.

TABLE OF CONTENTS



FROM TERRY WALLACE

PRINCIPAL ASSOCIATE DIRECTOR FOR SCIENCE, TECHNOLOGY, AND ENGINEERING

Information Science and Technology Revolution

INSIDE FRONT COVER



FEATURES

Roadrunner

COMPUTING IN THE FAST LANE

2



Magnetic Fields of Dreams

THE WORLD'S HIGHEST FIELDS UNVEIL EXOTIC MATERIAL BEHAVIOR

8



Close Encounters of the Particle Kind

THE MONTE CARLO CODE MCNP TACKLES ALL THINGS NUCLEAR

15



DIALOGUE

Agent-Based Computer Models

A COOL TOOL FOR POLICY PLANNING

20



SPOTLIGHT

BETTER BREAST CANCER DETECTION AND DIAGNOSIS

GETTING INSIDE A FLY'S HEAD

THE SOUND OF ONE STAR FALLING

EXPLOITS WITH SUPERBRIGHT LIGHT

24



Roadrunner

Computing in the Fast Lane

The drive for more computing power is running into a brick wall. Old speed-up tricks no longer work, and a new paradigm is needed to keep up with Moore's law, a well-known trend in computing. Los Alamos' Roadrunner supercomputer is blazing that new trail, using a video game chip to accelerate supercomputing performance.

Facing page, left to right: Sriram Swaminarayan, Ben Bergen, John Turner, Mike Lang, Tim Kelley, and Jamaludin Mohd-Yusof, represent the programming teams that proved Roadrunner could achieve its performance goals.

For the last two decades, the number of transistors on a computer chip has doubled every 18 months, a phenomenon known as Moore's law. The net thousand-fold increase has caused a parallel increase in computer performance, fueling a technological revolution in which ever-more-sophisticated electronic devices are transforming both economies and cultures.

The computer industry hopes to keep up with Moore's law, given the market incentives. The private sector wants greater access to information and multimedia entertainment—video learning tools, movies, video games.

Scientists want faster, more-powerful high-performance supercomputers to simulate complex physical, biological, and socioeconomic systems with greater realism and predictive power. The world's fastest supercomputer is Blue Gene/L at Livermore National Laboratory. It performs at nearly 500 trillion arithmetic operations a second (500 teraflop/s).

Los Alamos scientists are shooting for double that performance, or one petaflop/s, which is 1,000 trillion calculations per second. They plan to do it with the Roadrunner supercomputer scheduled for installation at Los Alamos starting this summer, with full operation targeted for early 2009.

Hitting that target is no small task; the building blocks of supercomputers—chips (microprocessors)—have reached their speed limits, endangering the future of Moore's law for scientific computing.

Screeching to a Halt?


A microprocessor depends on two things for its performance. Its internal clock speed determines how fast it does the number crunching—the logical and arithmetic operations. The number of transistors that make up its logic circuits determines how “smart” it is, that is, how many different types of operations it can perform simultaneously.

Modern supercomputers have thousands of identical computer nodes, each containing a microprocessor and

a separate memory. The nodes are connected to form a cluster and work simultaneously (in parallel) on a single problem.

Increases in supercomputer performance have come in part from increases in the number of nodes in the cluster but, more important, from increases in microprocessor clock speed. As transistors were reduced in size and placed closer together, it was possible to turn them on and off faster and thus increase the clock speed. Today's transistors in high-end microprocessors have shrunk to 65 nanometers (billionths of a meter) and are running at gigahertz clock speeds (billions of clock ticks per second).

Transistors are still shrinking, and the number per microprocessor is still growing, but performance measured in arithmetic operations per second has flattened out since 2002 (see below).



Moore's law, the doubling of transistors on a chip every 18 months (top curve), continues unabated. Three other metrics impacting computer performance (bottom three curves) have leveled off since 2002: the maximum electric power a microprocessor can use, the clock speed, and performance (operations) per clock tick.



Alex and Andrew Turner playing a game of “MX vs. ATV Untamed” on the Cell-powered Sony PlayStation 3. A similar chip will power Roadrunner.

“The biggest reason for the leveling off is the heat dissipation problem,” says Ken Koch, one of the leaders of the Roadrunner project. “With transistors at 65-nanometer sizes, the heating rate would increase 8 times whenever the clock speed was doubled, outstripping our ability to carry the heat away by standard air cooling.” It would be like running a car at

high speed with no water in the radiator.



Another huge obstacle to increased performance is the memory barrier. In the not-too-distant past, the time to fetch data from the node memory and load it into the processing units (called the “compute core”) of a microprocessor was comparable to the time it would take that core to do the number crunching. Now the number crunching is 50 times faster than the time to fetch and load data. The time spent in data retrieval and communications can no longer be ignored.

Clearly the old solution for increasing supercomputer performance—miniaturizing circuits and using faster clocks—is breaking down.


Video Games Open a New Path

“We replace our high-performance supercomputers every 4 or 5 years,” says Andy White, longtime leader of supercomputer development at Los Alamos. “They become outdated in terms of speed, and the maintenance costs and failure rates get too high.”

With this turnover rate, you might think the Lab would be a major force in computer development. “That used to be the case,” says White, “but today we’re small potatoes. Even a decade ago, when the Department of Energy purchased Blue Mountain for Los Alamos and Blue Pacific for Livermore at a total cost of roughly \$200 million, a German bank was spending much more, signing an \$8-billion agreement for information technology services with IBM. And the personal



The Cell microprocessor contains a Power PC compute core that oversees all the system operations and a set of eight simple processing elements, known as SPEs, that are optimized for both image processing and arithmetic operations at the heart of numerical simulations. Each is specialized to work on multiple data items at a time (a process called vector processing, or SIMD), which is very efficient for repetitive mathematical operations on well-defined groups of data.





computer market has exploded since then, bringing prices for commodity microprocessors way down. Since the late 1990s, because of cost, we've had little choice but to build supercomputers from off-the-shelf rather than custom-made components."

That's why in 2002, when Los Alamos scientists were planning for their next-generation supercomputer, they looked at the commodity market for a way to make an end run around the speed and memory barriers looming in the future.

What they found was a joint project by Sony Computer Entertainment, Toshiba, and IBM to develop a specialized microprocessor that could revolutionize computer games and consumer electronics, as well as scientific computing.

These corporations invested \$400 million dollars over 4 years to produce the Cell Broadband Engine (the "Cell"), a powerhouse microprocessor carrying 240 million transistors. Its first application would be in the Sony PlayStation 3, a top-end video-game console.

Today's video games are like interactive movies, complete with elaborate, computer-generated backgrounds and interacting characters that are more and more realistic.

The Cell was designed with enough computer power to enhance interactivity, allowing video games to be even less scripted. It has eight specialized processing elements (SPEs) that get around the speed barrier by working together. They can generate dynamic image sequences in record time, sequences that reflect the

game player's intention and even have the correct physics.

The Cell gets around the memory barrier as well. It does so by having a small, fast local (on-chip) memory plus a memory engine for each SPE and an ultra high speed bus to move data within the Cell. The local memories store exactly the data and instructions needed to perform the next computations while all eight memory engines act like runners, simultaneously retrieving from off-chip memory the data that will be needed for computations further down the line.

Optimized for maximum computation per watt of electricity, the Cell looked like a good bet for accelerating supercomputing performance. Los Alamos knew, however, that the Cell would need some modifications for petaflop/s scientific computing. IBM was willing to work on the enhancements.

A Hybrid That Raises Skepticism

Named after the fleet-of-foot New Mexico state bird, the Roadrunner supercomputer is a hybrid, containing not one type of microprocessor but two.


Its main structure is a standard cluster of microprocessors (in this case AMD Opteron dual-core microprocessors). Nothing new here except that each chip has two compute cores instead of one. The hybrid element enters the picture when each Opteron core is internally attached to another type of chip, the enhanced Cell (the PowerXCell 8i), which has been designed specially for Roadrunner. The enhanced


Cell can act like a turbocharger, potentially boosting the performance up to 25 times over that of an Opteron compute core alone.

The rub is that achieving a good speedup (from 4 to 10 times) is not automatic. It comes about *only* if the programmers can get all the Cell and Opteron microprocessors and their memories working together efficiently.

"The typical challenge in programming a code is to solve the mathematical equations in the fewest arithmetic (data processing) operations. Here the challenge was different—to minimize the flow of data while simultaneously taking full advantage of the computational power of the Cell," explains Sriram Swaminarayan, a Los Alamos computational physicist. "At first most of

Roadrunner is a cluster of approximately 3,250 compute nodes interconnected by an off-the-shelf parallel-computing network. Each compute node consists of two AMD Opteron dual-core microprocessors, with each of the Opteron cores internally attached to one of four enhanced Cell microprocessors. This enhanced Cell does double-precision arithmetic faster and can access more memory than can the original Cell in a PlayStation 3. The entire machine will have almost 13,000 Cells and half as many dual-core Optérons.





us were skeptical it could be done.”

Los Alamos had just over 12 months after signing the original contract with IBM, until October 2007, to decide whether it would purchase the Cells and proceed with its investment in this new hybrid computer architecture.

“Before making that decision, we needed to see if the Cell’s stringent programming demands could be dealt with successfully on the types of codes that are important to nuclear weapons stockpile stewardship and other national security missions,” comments John Turner, leader of the Roadrunner Algorithms and Applications team.

The Programming Experience

“The Cell is designed to move data in, compute on it, and move it out faster than an ordinary microprocessor but *only* if the programmer writes code in a specific way,” explains Koch. For the Cell, the programmer must know exactly what’s needed to do one computation and then specify that the necessary instructions and data for that one computation are fetched from the Cell’s off-chip memory in a single step. They are then stored in the on-chip memories of each of the Cell’s eight SPEs. IBM’s Peter Hoftstee, the Cell’s chief architect, describes this process as “a shopping list approach,” likening off-chip memory to Home Depot. You save time if you get all the supplies in one trip, rather than making multiple trips for each piece just when you need it.

The small size of the on-chip memories is an additional challenge. The programmer must divide the computation (the data and instructions) into chunks appropriate for the on-chip memories, then feed the Cell

Above: Recent VPIC simulations on Opteron-Cell nodes reveal that magnetic reconnection (energy transfer from magnetic fields to plasma particles) involves “kinking” the current layer and forming rope-like structures. Right: If chosen to run on Roadrunner, supernovae calculations using Milagro will be the first to determine the real influence of radiation flow on the light signals from these exploding stars.

many small chunks in an assembly line fashion; otherwise, the Cell will not cause a speedup.


Turner organized teams of physicists and computer scientists to restructure codes for a spectrum of important application areas, re-implementing sections as necessary for the new architecture. The teams tested the rewritten codes first on a single Opteron compute core plus one Cell microprocessor and then in parallel on up to 24 such Opteron-Cell pairs.

The major application areas addressed were radiation transport (how radiation deposits energy in and moves through matter), neutron transport (how neutrons move through matter), molecular dynamics (how matter responds at the molecular level to shock waves and other extreme conditions), fluid turbulence, and the behavior of plasmas (ionized gases) in relation to fusion experiments at the National Ignition Facility at Livermore National Laboratory. The corresponding codes represented a range of methods for solving equations on a computer.

In the end, each code achieved a substantial speedup when run on a Cell-accelerated Opteron compute node in comparison with execution on a single Opteron compute core, without the Cell.

The VPIC code, which simulates plasmas in magnetic fields, is a prime example. It ran 6 times faster on the Opteron-Cell node than on the Opteron alone. That increase will allow researchers to tackle some scientific grand-challenge problems.

Successfully accelerating the Monte Carlo code called Milagro took many months, several false starts, and modification of 10 to 30 percent of the code. Monte Carlo codes, which simulate radiation transport, are very expensive computationally. As the October decision time drew near, Milagro was also executing



6 times faster with the Cell than without, a crucial achievement for the acceptance of Roadrunner.

Bringing It All Together


But testing the codes on single nodes or small groups of them could provide only part of the performance picture. “The performance of systems such as Roadrunner comes from a complex interplay of system architecture and the workload characteristics of the applications,” says Adolfo Hoisie, who leads system and applications performance activities for the Roadrunner project.

To project how the codes would perform on the full system, with its average of 3,250 compute nodes containing about 13,000 Cells, computer scientists in the Performance and Architecture Lab (PAL) brought out all their novel and proven performance-modeling methodology.

PAL’s work was instrumental not only in accurately indicating the potential performance advantages of Roadrunner compared with other architectures but also in quantifying and guiding various system design decisions that ensured top performance on the codes of greatest interest to the Lab.

By the time October 2007 rolled around, Los Alamos was confident that the entire range of science problems, from radiation transport to molecular dynamics, would run on Roadrunner at accelerated speeds, from 4 to 9 times faster than on the Opteron cluster alone.

Los Alamos scientists are now confident that Roadrunner will become the world’s fastest supercomputer. It will be a tremendous asset to the computer simulations performed at the Laboratory for the nuclear weapons program as well as for scientific grand challenges. Important codes are expected to run at 200 to 500 teraflop/s. Roadrunner will also be the first



High-resolution simulations (30 billion cells) of Rayleigh-Taylor turbulence reveal details of the mixing layer between two fluids.

computer to run the universally recognized code used to test supercomputer performance—LINPACK—at over 1 petaflop/s.

Hybrid Computing—The Wave of the Future

After the late-summer delivery and a check-out period, Roadrunner will be opened to unclassified science applications for 4 to 6 months. A call for proposals within Los Alamos has already been issued. Turner says,

“We expect to see proposals in cosmology, antibiotic drug design, HIV vaccine development, astrophysics, ocean or climate modeling, turbulence, and we hope many others.”

Afterward, Roadrunner will be moved to the classified computing network and used to improve specific physics models for the nuclear weapons program and for validating the accuracy of answers from earlier, less-detailed nuclear weapons simulations.


By 2010 the Lab’s scientists plan to use Roadrunner to help definitively quantify uncertainties in simulations of nuclear weapons performance as well as to reduce those uncertainties in key areas. This is an important milestone in maintaining confidence in the nation’s nuclear weapons stockpile without actual nuclear testing.


In the meantime, Roadrunner is seen as the first in a new wave of hybrid supercomputers that will be used for scientific grand challenges around the world.

Turner sums up the Los Alamos outlook this way, “We broke new ground in learning to program Roadrunner, but the people who did it found it fun. Our task now is to transfer the knowledge we gained so that new users of hybrid supercomputers won’t have to go down the blind alleys we explored.” ❖



Above: Cooling towers at the Los Alamos Metropolis Center for Modeling and Simulation dissipate the heat generated by the power-hungry supercomputers.





Magnetic Fields of Dreams

Scientists use strong magnetic fields to investigate the inner workings of new, often-exotic materials. With the ability to create some of the highest magnetic fields in the world, Los Alamos is playing an important role in helping to understand the materials of tomorrow.

With the flick of a few computer-controlled switches, a giant pulse of electricity races through the powerful electromagnet at the Los Alamos Pulsed Field Facility (the Magnet Lab). An immense magnetic field builds within the magnet's center, a field more than a million times stronger than Earth's and having the stored-energy equivalent of a few sticks of dynamite. Physicist Ross McDonald wonders whether it is strong enough to allow him to study the enigmatic electrons of a high-temperature superconductor.

Scientists at the Magnet Lab use strong ("high") magnetic fields to gain a fundamental understanding of a broad range of materials, from conducting metals to nonconducting insulators, from stringy polymers to heat-resistant ceramics. Often, these are new materials, hot off the laboratory bench and of interest to scientists because of what's going on inside them.

Consider the inner workings of metals. At the atomic scale, metals consist of positively charged atoms (ions) held together in a three-dimensional lattice by negatively charged electrons. Most of those

electrons are bound to the ions, but some are free to roam through the lattice and, for example, conduct electricity.

In many of the metals brought to the Magnet Lab, those free electrons behave strangely. In some, the electrons act as if they were obese particles with masses hundreds of times greater than those of normal electrons. In others, the free electrons, which normally avoid each other, cluster at regular positions within the material, creating an unusual secondary lattice within the main one—a "charge-density wave."

Then there are superconductors. When these metals are cooled close to absolute zero temperature (-273°C), the free electrons form "pairs," a special quantum state that allows both electrons to move through the lattice without resistance. As a result, electricity flows through a superconductor without energy loss.

In ordinary superconductors, the electron pairs break apart even before the temperature rises above a frigid -250°C , and the metal becomes a normal conductor of electricity. That's not the case for the high-temperature




Ross McDonald threads a probe through the top of the 100-tesla multi-shot electromagnet, which extends downward about 7 feet. The large magnet is one of a half-dozen high-field magnets at the Magnet Lab.

superconductor being studied by a Los Alamos team including John Singleton Neil Harrison, Chuck Mielke, Fedor Balakirev and Ross McDonald. The superconductor's electrons remain paired for another 60 degrees or so. And researchers can't agree on why.

The scientists who study heavy-electron behavior, charge-density waves, or high-temperature superconductivity are the first to admit that the

particular metals they study have few applications and are unlikely to make for better bridges or niftier iPods. But they also know that an understanding of the arcane electron behaviors in those metals will give scientists a greater chance of creating the next-generation materials that will make our buildings greener, our gadgets smaller, our power and light systems more efficient—or whatever else we can imagine.



Left: Electric current sent through a wire coil creates a beautifully symmetric magnetic field that gets stronger as the current increases, the wire loops get smaller, or the number of loops increases. If the current is delivered as a pulse, a pulsed magnetic field is created.

Recently, some Magnet Lab staff took a major step toward deciphering the electron behavior in a high-temperature superconductor by measuring a distinctive and telling feature of the metal—its so-called Fermi surface. Such a measurement had eluded scientists for more than 20 years but was achieved using a cutting-edge measurement technique and the Lab's powerful new magnet, the "100-tesla multi-shot."

Higher Fields Are Better

Some of the highest pulsed magnetic fields on the planet are created at the Los Alamos Magnet Lab. The facility is part of the National High Magnetic Field Laboratory (NHMFL), a research center with headquarters at Florida State University in Tallahassee and additional facilities at the University of Florida in Gainesville. A user facility funded principally by the National Science Foundation but also by the State of Florida and the U.S. Department of Energy, the NHMFL generally affords users free access to the magnets for research in chemistry,



biology, physics, geophysics, and medicine.

“Here at Los Alamos, the focus is on condensed-matter physics and materials science,” says Marcelo Jaime, interim director of the Los Alamos facility. “We’re also one of the few places that can conduct investigations into uranium or plutonium metals, as well as materials containing the two elements.”

As an investigative tool, a magnetic field is analogous to temperature and pressure; it’s something a scientist can control to change a material’s properties. By mapping out how a property changes with field strength and comparing the results with theory, scientists can begin to understand the more-unusual electron interactions.

The changes come in part because a magnetic field interacts with the electron’s spin—an intrinsic property that makes the electron act like a tiny magnet. The electron responds to the field like a subatomic compass needle, aligning its spin with or against the field direction. Depending on the alignment, the electron’s energy shifts up or down by a small amount that’s proportional to the field strength. The field also affects the path of a moving electron, pushing it sideways and altering its momentum.

“The field is like a big lever we use to induce significant changes in a material,” says McDonald. “With higher fields, we can align more spins and shift the electron’s energy or momentum enough to disrupt some of the stronger electron interactions, such as electron pairing in high-temperature superconductors. We couldn’t have measured the Fermi surface without a really high field.”

Electrons at the Fermi Surface

The number of electrons in most metals is astounding—in the neighborhood of a trillion trillion per cubic centimeter. Each of those electrons has a unique combination of energy, momentum, and spin alignment.

When a metal’s total energy is at its theoretical minimum (at a temperature of absolute zero), its electrons assume every available energy at that temperature, from the lowest value to some very-high value known as the Fermi energy. A three-dimensional plot of the Fermi energy as a function of the electron momentum is the Fermi surface.

Because all energies below the surface are taken, most of the metal’s electrons are locked in place energy-wise. They can’t get enough additional energy (through collisions or whatnot) to access one of the available energies lying above the surface, so they don’t respond to external influences. Only electrons on or close to the Fermi surface can change their energy and respond.

The Fermi surface is therefore a map of the metal’s “important” electrons, the ones responsible for almost all of its electronic properties, and is an invaluable “reality check” for theorists trying to predict those properties. The shape of the Fermi surface



Above: The Magnet Lab’s generator, the largest in the world, can supply second-long pulses of electricity to power the large magnets. Right: Transformers (red cylinders) shape the voltage pulse from the generator into customized pulses, producing fields that rise and fall as desired.

A Mighty Magnet

A magnetic field is created when electricity runs around the circular turns of a wire coil. It takes hundreds of thousands of ampere-turns (current times the number of turns) to produce a high field, say, above 25 teslas. (A common refrigerator magnet has a strength of about 0.1 tesla.)

A large coil may have hundreds of turns, and the thousands of amperes needed to produce 25 teslas generate enough heat (in a resistive, non-superconducting magnet) to melt the coil windings. That same field exerts a pressure on a piece of reinforcing steel that is more than twice what's felt at the bottom of the ocean. Designing an electromagnet that can survive a multi-tesla field and be used over and over again is remarkably hard.

"Every high-field magnet will eventually break down because we push the limits of what the materials can withstand," says Chuck Swenson, leader of the Pulsed Magnet Design project. "The challenge in producing ever-higher fields is to understand electromagnets well enough to create new, durable designs."

Those designs are works of functional art, highly optimized and consisting of multiple coils placed one inside the other like a set of Russian dolls. Wherever and whenever possible, coils are wound from copper impregnated with nano-scale ribbons of niobium—an extremely strong yet excellent conductor. Reinforcing steel laced between the coils helps maintain their structural integrity.

All of the highest-field magnets are pulsed: a single swift current pulse sent through the assembly creates a field that rises, peaks, and decays (typically) within a few thousandths of a second. The short duration limits the heat and stress on the materials, so the highest fields can be contained without destroying the magnet.

More than a decade of research has culminated in the

determines the electrons' movement through the lattice, the metal's optical properties, and the likelihood that the metal will alter its electronic and/or crystal structure when subjected to stress. And by looking at how the measured surface differs from its predicted shape, scientists can infer new interactions that could explain strange electron behaviors.

A Super Conductivity Measurement

The high-temperature superconductors are metallic compounds, also known as cuprates for the copper that is part of their composition. Since the cuprates'



Top: The Large Magnet Assembly Team surrounds the 100-tesla multi-shot magnet, seen without its central insert coil.
Bottom: The insert coil being loaded into the magnet.

Magnet Lab's 100-tesla multi-shot, currently the most-powerful reusable magnet in the world. Designed to operate at 100 teslas, the multi-shot has so far been kept to 89.9 teslas, still a world record.

The magnet's commissioning late in 2006 added significant capabilities to the Magnet Lab, including the ability to break the superconducting state of a frozen cuprate.

discovery in 1986, numerous theories have been proposed to explain why they remain superconducting at higher-than-usual temperatures. There was even speculation that these materials were not metals in the usual sense but were conducting electricity through some unknown mechanism. That's because all attempts to measure a cuprate's Fermi surface failed to prove it existed at all.

The unmistakable signature of the Fermi surface is the oscillation of the value of some property, such as the conductivity (a measure of the material's ability to conduct electricity), as an applied magnetic field

increases. Such oscillations occur because the field shifts the energy of the electrons. At some field values, lots of electrons move onto the Fermi surface and the conductivity increases, while at others, the electrons move off the surface and the conductivity decreases.

“The oscillations are small,” says Singleton. “Typically, there are three requirements needed to observe them: very-pure single-crystal samples, very-low temperatures, and a very-high magnetic field. For the cuprates, all three were problematic.”

First, the metals are far from being pure crystals. They’re made by adding atoms (dopants) to the material in a disorderly manner that wreaks havoc on the crystal structure.

The low temperature was a second problem. Heat makes the ions vibrate, the vibrating ions bump electrons from their paths, and signals produced by the electrons become noisy. The noise can be reduced by making the measurements at very-low temperatures, but then the free electrons in the cuprates pair up and become superconducting. Unfortunately, pairing lowers the electron energy, so all the electrons leave the Fermi surface, making it impossible to measure its shape.


The solution to this second problem was to apply an immense magnetic field that would shift the electron energy enough to break the superconducting pairs apart while the metal was at low temperatures. The electrons would remain on the Fermi surface, and science could move forward.

Producing that high field was a third problem. The Magnet Lab’s solution was to develop the 100-tesla multi-shot—one of the most-powerful magnets in the world. (See “A Mighty Magnet.”)

The Race Is On!

By 2007, more than two decades after high-temperature superconductivity was discovered, all the experimental obstacles had been overcome. Crystal growers had developed a technique to grow a cuprate in which the dopants were arranged in an orderly lattice. They were able to hand the experimentalists what amounted to a pure crystal. In addition, magnets of the necessary strength had been built.



Remarkably, the Los Alamos team and a Canadian group were racing to be the first to measure the Fermi surface of a cuprate metal. The Canadian group was hunkered down at the Laboratoire National des



The Fermi surface reveals how the energy varies with momentum for the highest-energy electrons—those that have the Fermi energy. Left: The simple Fermi surface of copper. The blue outline relates to the lattice. Every electron with a momentum that lies on the yellow surface has the Fermi energy. Below: The strange Fermi surface (blue tubes) of a high-temperature superconductor. The inner box relates to the lattice. This surface is consistent with a model in which the spins of neighboring ions point in opposite directions.

Champs Magnétiques Pulsés in Toulouse, France. The Los Alamos team was using the 100-tesla magnet and an advanced technique, contactless conductivity, to make the measurement.

“We spent a lot of time developing and refining contactless conductivity,” comments Mielke, head of the Magnet Lab’s user program. “A tiny electrical coil surrounds, but doesn’t touch, the sample. The coil and sample each have an inductance, or magnetic resistance. When the field is off, the inductance causes a tuned circuit to resonate at a certain frequency. When the field turns on



Prototype insert coil for the 100-tesla multi-shot. Consisting of 16 coils layered in two sets of 8, this assembly should allow the multi-shot to reliably produce fields of 100 teslas.



and increases, the sample's conductivity oscillates in value, which causes its inductance—and its resonance frequency—to oscillate as well. From those frequency oscillations we can deduce, with amazing precision and resolution, the shape of the Fermi surface.”

Contactless conductivity did the trick. The Los Alamos team succeeded in its measurement, but not before the Canadian researchers had succeeded in theirs.

Now that the Fermi surface is known to exist, it's clear that traditional theories of metals do indeed apply to the cuprates. “The question remains as to whether the surface will yield enough information for us to tease apart the electron interactions that govern the pairing mechanism,” says Singleton.

New Frontiers

The Los Alamos team is currently investigating how the Fermi surface evolves as the cuprate's composition changes. In comparing all the data (including the controversial results from an experiment conducted at Los Alamos in 1991), one sees dramatic changes in the Fermi surface as the materials get closer to the number of dopants that is optimum for the highest superconducting temperature.

The Magnet Lab is continuing its quest to produce higher fields. Indeed, Mielke is spearheading a new

electromagnet design, the “single-turn,” named for its single loop of copper. The single-turn has already produced pulsed fields as high as 240 teslas. The field lasts but a few millionths of a second, and then—the magnet explodes! Remarkably, the magnet's design allows a sample to survive the explosion intact.

Mielke is planning to use the single-turn to measure the Fermi surface of plutonium and to investigate superconductivity in the heavy-electron metals, but he needs to refine his measurement techniques. “A changing magnetic field can generate an unwanted voltage—electromagnetic interference (EMI)—in the measurement probe,” he explains. “It's hard enough to measure small signals in the 100-tesla magnet, where the field goes from nothing to everything in a few thousandths of a second. When the field ramps up in the single-turn's millionth of a second, the EMI is much higher, and the measurement becomes that much harder.”

Mielke is patiently refining his techniques, the same way that all Magnet Lab scientists refine and advance theirs. They all recognize that tomorrow's materials will likely be discovered through an understanding of today's and that gaining such understanding is a slow process. But high magnetic fields are the quickest way. ❖

Above: Chuck Mielke with the single-turn electromagnet. The white tube passing through the magnet's single copper loop contains the sample and the measurement probes. A 3-million-ampere current pulse creates a 240-tesla magnetic field at the loop's center, generating so much heat and magnetic force that the loop explodes. The tube and its contents survive intact. Inset: A successful experiment.



John Hendricks uses dice to help teach prospective MCNP users how the code works.

CLOSE ENCOUNTERS OF THE PARTICLE KIND

With many years of science built into it, the Los Alamos Monte Carlo code called MCNP can tell you how subatomic particles will interact with just about anything.

John Hendricks lets his students play with dice. Most of us associate the little cubes with gambling, but for him they're a great teaching tool.

Hendricks, a member of Los Alamos National Laboratory's Applied Physics Division, teaches people to use what he says "is one of the Lab's greatest achievements."

No, it's not "the bomb." It's MCNP—Monte Carlo N-Particle—a computer code that has indirectly touched all of our lives. The oil in our cars was likely found with help from engineers using MCNP (or its predecessor versions, or MCNPX, all of which are referred to here as MCNP). The air we breathe was possibly monitored by a system that relies on MCNP, while anyone who has undergone radiation treatment received a dose that was likely calculated and verified by the code.

The uncommon link between these widespread applications is that each requires knowing what happens to radiation—free particles such as neutrons and photons—as it makes its way through matter. Modeling the "radiation transport" of gamma-ray photons through the body can tell you how many will be absorbed by cancer cells, while following the transport of neutrons from a source to a detector, both placed underground, can tell you if the neutrons passed through oil-saturated rock.

Computationally, it turns out that the most accurate way to solve a radiation transport problem is to use random numbers to decide the outcome of a particle's encounters with atoms and nuclei, and repeat the process for many particles. It's a physics-based game of chance, and MCNP is the game's most-skilled player.

A Game of Chance

In the classroom, Hendricks poses a question to his students. "Suppose you want to design a nuclear

The heavenly bodies of our solar system are constantly bombarded by cosmic rays, causing elements on their surface to emit radiation. By using MCNP to simulate those emissions and their detection by a satellite system, scientists were able to identify water on Mars.

reactor. How do you figure out whether your design produces energy safely, reliably, and sufficiently? In part, you do it by using MCNP to keep tabs on the neutrons."


When a neutron is absorbed by the oversized nucleus of a uranium atom, it can cause the bit of matter to fission, a reaction whereby the nucleus splits into two pieces (fission fragments). Fission unleashes a relatively huge amount of energy—millions of times more than is released from a chemical reaction between atoms.

Significantly, a few neutrons are also unleashed in the nuclear breakup. In what's known as a chain reaction, those freed neutrons can cause more nuclei to fission, unleashing more energy and even more neutrons, and so forth. A nuclear reactor's primary job is to achieve a steady state chain reaction: to keep fission going and the energy flowing.

But other kinds of nuclear reactions may occur when a neutron meets up with a uranium nucleus.



CREDIT: NASA AND STSCI



CREDIT: FRENCH NUCLEAR NEWSLETTER.

For example, the neutron may be captured—the uranium absorbs it but does *not* fission. Or the neutron may simply bounce off a nucleus and scatter. One cannot determine in advance what happens in any specific encounter, but can only determine the *probability* for a given reaction to occur.

MCNP knows the different probabilities from experimental data. It then uses one random number to choose between reactions (fission, capture, scattering, or other), and others to select a specific outcome: the number of neutrons and/or other particles produced, their final energies, the direction each taken by each, even how far each goes before encountering a new nucleus. The program generates its random numbers with a simple algorithm. Hendricks and his students

obtain them by tossing dice. (See “A Roll of the Dice”)

Making independent choices for each new encounter, MCNP is able to construct a realistic trajectory of a neutron through the reactor. It builds such a trajectory for thousands, if not millions, of neutrons, enough to create a statistically accurate picture of their fate and consequences. MCNP can then calculate the rate at which neutrons are produced by fission and the rate at which they are lost. When the production rate equals the loss rate, you get a steady-state chain reaction, which is the beating heart of a nuclear reactor.

The Best of the Best


In many ways, MCNP, with its attendant nuclear data sets, is the repository of all we know about the interactions between radiation and matter. It’s not the only code that deals with those interactions; it’s one of many “Monte Carlo” codes, so called for their reliance on chance (see “The Code Was in the Cards”). But if it’s not unique, it *is* the gold standard, its superiority anchored by a complete set of physics data, efficient ways to reduce statistical uncertainty, a sophisticated graphics interface, and an uncommonly bug-free source code.

“It’s the best of the best,” says Hendricks, who, like Art Forster, Tom Booth, and Gregg McKinney of the Applied Physics Division and Laurie Waters of Decision Applications Division, has devoted his career to MCNP. “It accurately describes anything and everything that can happen in a particle-nucleus collision.”

The Advanced Fuel Cycle

MCNP is being used by nuclear engineers and scientists, including some at Los Alamos, to address one of the more pressing issues of our time—the production of nuclear energy. The goal is to develop new reactors, new nuclear fuels, and enhanced safeguard systems that will make nuclear energy safer, produce less waste, and be resistant to the diversion of plutonium.

In much of the world, nuclear power is produced using an “open” fuel cycle: enriched uranium fuel is burned (fissioned) in a reactor until it can no longer produce energy economically, at which point the fuel is deemed spent, removed from the reactor, and considered waste to be discarded. But spent fuel still has lots of energy-rich



Above: A French advanced fuel assembly. Right: This figure shows MCNP predictions for the radiation signals emanating from a person contaminated with a known dose of the radioactive isotope cobalt-60. Note that cobalt-60 accumulates primarily in the lungs and liver.

fissionable material, so the open cycle is analogous to a business that doesn't recycle.

The scientific understanding needed to close the fuel cycle and implement a recycle strategy already exists. The U.S. Department of Energy's Advanced Fuel Cycle Initiative (AFCI), which is the research and development arm of the Global Nuclear Energy Partnership initiative, seeks to optimize the more mature technologies and develop new ones that would enhance energy extraction from the nuclear fuel, minimize waste, and reduce proliferation risks. This includes developing and refining technologies to chemically reprocess the spent fuel, that is, remove the uranium and materials known as transuranics, and turn them into an entirely new type of fuel that can be burned for further energy production.

The transuranics, elements heavier than uranium (neptunium, plutonium, americium, etc.), are created in the fuel by nuclear processes. Plutonium, for example, is created after uranium captures a neutron and decays to neptunium, which then decays to plutonium.

Transuranics are the Methuselahs of spent fuel, taking hundreds of thousands of years to decay. That time scale underlies all efforts to stabilize the waste and develop repositories for its storage or disposition. By burning the transuranics, the advanced fuel cycle squeezes out more energy from each bit of fuel and creates a much simpler waste disposal process. Researchers are developing optimized mixtures of uranium, plutonium, and other transuranics for the new TRU fuels (so-called for their enhanced TRansUranic component).

One complication is that today's most-common commercial reactors, thermal reactors, don't burn TRU fuel efficiently. Thermal reactors depend on slow, or low-energy, neutrons for fission. The new fuels would require the use of fast-neutron reactors, designed to sustain fission with neutrons of very high energy. Through the AFCI Program, DOE is supporting research into this next generation of reactors, and MCNP is being used to help design them.

Meanwhile, work on the new fuels is underway, and once again, MCNP is involved.

A Roll of the Dice

How can choosing a random number lead to physically relevant descriptions? John Hendricks explains how to a class of young engineers:

"Suppose that in an imaginary nuclear reactor, there are only four possible outcomes for a neutron-nucleus encounter: (A) The neutron is absorbed and captured by the nucleus, so no neutrons remain after the encounter; (B) the neutron scatters off the nucleus, so one neutron remains; (C) the neutron is absorbed and the nucleus fissions, producing two neutrons, or (D) the nucleus fissions and produces three neutrons.

"Next, imagine that experiments with our nucleus reveal that for every six neutron encounters, (A) happens 3 times, and (B), (C), and (D) each happen once. How do we model this physics?"

"We'll be rolling a die that generates a number from 1 to 6 at random and with equal probability. We have to make the six numbers correspond to the four possible outcomes in a way that matches the probabilities that were measured in our experiments. This is done by assigning the die numbers 1, 2, and 3 to (A), 4 to (B), 5 to (C) and 6 to (D). In essence, we've transformed the 'real world' into an equivalent 'computational world,' where each encounter now has six outcomes that result in either 0, 0, 0, 1, 2, or 3 neutrons. Because our die is unbiased, each outcome is equally likely to occur, so on average, six rolls of the die will produce a total of six neutrons.

"To find out how our reactor behaves, we do a Monte Carlo simulation. We'll start with many neutrons in the reactor. Each neutron will encounter a nucleus, and a roll of the die will determine the outcome (how many neutrons remain). We'll then average our results and calculate the uncertainty of our answer.

"If we run our simulation with a million neutrons, we will find, on average, that a million neutrons remain, that is, our reactor achieves a steady-state chain reaction! The relative uncertainty of our result will be about 0.1 percent. One of MCNP's most powerful features is its use of mathematical shortcuts that reduce the uncertainty without having to 'roll dice' a million times. Thank goodness for that!"





Left: Anna Hayes of Theoretical Division heads an effort to merge MCNP with other programs to create a reactor-performance program.

Bottom: Nuclear Nonproliferation's Martyn Swinhoe holds a neutron-sensitive detector, one of many stationed vertically inside the rectangular frame beneath his hands. Fuel rods containing fresh uranium/plutonium fuel are lowered through the frame. Neutrons produced by the plutonium are detected. MCNP interprets the neutron signal and verifies the stated plutonium content.

MCNP and the Materials Test Station

The road to the advanced fuel cycle will pass through a new facility planned for construction at the Los Alamos Neutron Science Center (LANSCE): the Materials Test Station (MTS). Expected to open in 2012, the MTS will provide a steady stream of neutrons that approach, both in energy and in flux (neutrons per square centimeter per second) the flow expected within a fast-neutron reactor. Candidate TRU fuels for the new reactors, some of which are being developed at Los Alamos, can thus be tested in advance.

The engine powering the MTS is the LANSCE facility's 800-million-electronvolt proton beam. When the protons slam into the MTS's tungsten target, copious neutrons will be produced (about 20 per proton-tungsten nucleus collision). MCNP is critical for simulating the results of those collisions and for following the torrent of neutrons as they zip and scatter their way into the candidate fuels.

"We'll measure how much of the fuel's transuranic material fissioned during irradiation," says Eric Pitcher, the test station project manager. "MCNP will tell us the neutron flux, so we'll know how effectively each fuel burns. We'll also monitor a host of other things, including how much radiation damage is done to the fuel rod material (the cladding). The cladding maintains an all-important barrier between the fuel and the coolant and must perform with a high degree of integrity and robustness."

One Code to Model Them All

While MCNP can tackle any radiation transport problem, it can't model the bulk properties of the


reactor materials or predict other types of behaviors. For example, the circulation of the liquid coolant in a reactor needs to be modeled by a "hydrodynamics" code.

Anna Hayes and her colleagues in Theoretical Division are developing an "omnibus" design code that combines MCNP with a hydrodynamics code and a nuclear fuel burn-up code. It is also tied to a code that calculates the equation of state (EOS) of a material, which tells you how the material's properties change as a function of temperature and pressure.


"It's important to track the flow of energy in a fast reactor very accurately. Fission fragments don't go very far, so they dump their kinetic energy within the fuel rod," says Hayes. "We estimate that energy with a fuel burn-up code, CINDER. The neutrons travel farther and tend to transport energy from one fuel rod to another. MCNP tracks the neutrons' energy and tells us the amount dumped into the cladding. The information is handed



Proton Flux



Neutron Flux



off to the ‘hydro’ and EOS codes to determine the impact of this heat on the cladding and coolant.”

One concern that the group will address is the production of helium and xenon gas bubbles in the fuel and cladding. They’ll be looking to see if the gases could affect the integrity of the fuel rod.

Keepin’ Things Safe

A final area where MCNP is playing a large role is in nuclear safeguards, that is, keeping track of plutonium and other fissionable materials throughout the nuclear fuel cycle. The Los Alamos Nuclear Nonproliferation Division’s Martyn Swinhoe uses the code to design radiation detectors for organizations such as the International Atomic Energy Agency that keep tabs on nuclear material worldwide. In the last year Swinhoe and his team have used MCNP to develop nine different monitoring instruments for four different organizations intent on surveying the nuclear material in the fuel reprocessing plant at Rokkasho, Japan.

MCNP determines the behavior of the neutrons that are emitted from plutonium in a fuel rod and reach the neutron-sensitive elements within the detector. The code provides Swinhoe with an estimate of the neutron detection rate (the total counting rate) and also the rate at which pairs of neutrons from a fissioning nucleus will be detected (the coincidence counting rate). The latter gives direct information on the mass of plutonium in the fuel rod.

Above: The Materials Test Station will verify the performance of advanced reactor materials and fuels. The top simulation, derived from MCNP data, shows two proton beams (colors) being stopped in the twin tungsten targets. Red is highest proton flux, blue is lowest. In the bottom simulation, fast neutrons produced by the protons permeate the fuels and materials placed within the circles and within the small rectangles surrounding the targets. Red is highest neutron flux, blue is lowest.

The Code Was in the Cards

In 1946, a game of canfield solitaire led Stanislaw Ulam, the Polish mathematician and Manhattan Project pioneer, to a brainstorm. Having returned to Los Alamos, Ulam was temporarily idled by illness and passed the time playing the game. He also pondered the odds of winning any given hand but was stymied by the combinatorial equations needed to exactly solve the probability. Instead, he hit on the idea of using statistical sampling: he would zero in on the probability of a successful outcome by laying out a large number of hands and keeping track of the winners.

Statistical sampling was not original to Ulam, but his musings on cards and chance led to two key insights. He first recognized that statistical sampling was well suited to the problem of neutron diffusion and next realized that the new electronic computers could generate the potential pathways of the neutron through the material. He shared these ideas with Hungarian mathematician John von Neumann, who at the time was a consultant to both Los Alamos and the Ballistics Research Laboratory in Maryland, where the first general-purpose digital computer, the ENIAC, was being built.



Liking the idea, von Neumann in 1947 laid out the neutron diffusion problem in a 19-step computing sheet: in essence, the first Monte Carlo code. The flamboyant name was given to the method by Nicholas Metropolis, architect of Los Alamos’ first computers. Ever since then, statistical-sampling computer codes have been known as Monte Carlo codes.


Stanislaw Ulam

“We use MCNP to optimize the design of the detectors before they are built,” says Swinhoe. “It saves us a lot of experimental work on prototypes.”

With nuclear energy making a comeback around the world, it’s not surprising to see MCNP turning up as the tool of choice for designing both the nuclear fuel cycle of the future and the safeguards to control where and how nuclear fuel and its waste products are used. ❖

Agent-Based Computer Models A Cool Tool for Policy Planning

Rob Axtell, Associate Director of the Center for Social Complexity at George Mason University, and Ed MacKerrow, a Los Alamos scientist and expert in agent-based modeling, discuss how computer experiments using interacting agents are becoming an important tool for policy development.



Ed MacKerrow

1663: Agent-based computer models have become a hot item in Washington policy circles. Can you explain what they are and why they're important?

Axtell: Agent-based computer models are a new way of working on social science issues. They are used to simulate complex social systems composed of interacting agents—either individuals (consumers, voters, commuters) or social groupings (families, firms, cities, nations).

MacKerrow: Basically, the simulations amount to science experiments in which we have full transparency into the rules of behavior and can control variables relative to each other. For example, a corporate manager might tell us what factors, including emotional ones, influence his decisions.

We then encode those rules of behavior into "if-then" algorithms in the agent-based model. By including rules of behavior for lots of different interacting agents and running the model many times, we can generate a group of virtual corporate histories. We then examine the interesting ones to discover how and why the agents behaved as they did in different situations.

Our goal is to capture the differences in behavioral drivers across a variety of people.

1663: This seems like a radical way to do social science.

Axtell: As a method, it's actually more empirically based than most. The social sciences have many armchair theories about why people do what they do, how institutions work, or why the stock market gyrates the way it does. But those theories are based mostly on speculation because experiments in the social sciences are hard to do.

MacKerrow: The agent-based approach actually prompts us to collect information about people's motivations and their personal incentives. Whereas personal incentives may be difficult to include in equation-based approaches, they are quite naturally represented in computational algorithms.

Axtell: We take the individual-level behaviors and try to figure out what's going to happen at the macroscopic level—the level of the company, the institution, or the whole society. Chris Langdon of the Santa Fe Institute described this modeling approach as a kind of "macroscope," as opposed to a microscope.

MacKerrow: Perhaps your model of retirement would be a good example to describe.

Axtell: The basic question is what causes people to retire. The standard economic view is that people are completely autonomous actors who make unilateral decisions based on their assets, age, current employment—basically their economic status.

What's missing is any sense of sociality, of wanting to conform to what your friends, neighbors, and co-workers are doing. So we built a model in which you don't feel like working anymore if all your buddies are retired and playing golf. And conversely, if your friends are all still working hard, then maybe you don't retire even though you'd like to.

Certain peculiarities in the historical retirement data seem more in line with our model than with the conventional models. For example, when retirement policy shifted in 1960 to permit people to retire at 62 but with reduced Social Security benefits, there wasn't an abrupt shift to the early retirement



age or a smooth steady shift downward. Instead, there was an irregular pattern of abrupt changes from year to year that settled down in the late '90s, with the most-common retirement age at around 62 instead of 65. Our model produces those kinds of abrupt, irregular changes as individuals try to coordinate with those around them.

MacKerrow: Rob calls it a cascading model. Interestingly, it's not unlike the way a nuclear fission reaction can cause a cascade. A few particles create more particles, and there's a sudden amplification process.

Axtell: Another example where agent-based models really work is in reproducing certain mystifying empirical regularities, like the distribution of city sizes and business firm sizes (in number of people), which have remained the same for hundreds of years.

The distributions don't follow the standard bell curve (Gaussian). If they did, the number of cities or firms that were bigger than average would decrease very rapidly. Instead, the number decreases more slowly

1663: Like a power law? The number of cities of a given size equals the size to a certain power?

Axtell: Exactly. There are many more large cities, firms, or what have you than a bell curve would predict.

MacKerrow: It's interesting that the variance around the average city or firm size is not well defined. For example, if you calculated an average U.S. firm size from measured data (there are 120 million workers and 6 million firms), the answer would be 20 people. In reality, there are an unexpectedly large number of firms with hundreds and even thousands of employees.

Axtell: That also means that large firms have a disproportionate influence on the economy compared with average-sized firms.

Agent-based models explain why the number of large cities is disproportionately high.

1663: So what kinds of behaviors do you have to allow in the models to get those power law distributions?

Axtell: Most economic theorists start with a model in which the system is in static equilibrium: people can't make changes that allow them to do better because they're already making all the possible tradeoffs. Those models can't produce the large fluctuations at the high end of the distribution that we see for firm and city sizes.

Because we see power law distributions across the social sciences, many of us think that social processes should *not* be modeled as equilibrium systems. They should be modeled as systems in a state of perpetual evolution or fluctuation in which people are constantly trying to better themselves and their living conditions—nothing is really stationary.

1663: What qualities allow people to better themselves? Do they have different skills and abilities? And do they fill in new niches, with opportunities expanding because of that diversity?

Axtell: When we talk about people bettering themselves and their situation, we're actually examining the difference between growth—simply getting bigger—and development. Whereas mathematics does not lend itself to building models with different qualities of growth, we can do that by giving the agents in our models some distribution of skills. And then we can ask: Which kinds of skills are present after a hundred years in the economy? Which kinds of skills have died out? How does technology evolve? The computational models can help us answer those questions.

1663: Are the agent-based methods being accepted?

MacKerrow: There's a pretty wide-scale acceptance across a variety of fields. The models give decision makers in Washington something tangible. It's not just someone's speculation.

Axtell: Plus, policymakers have a big appetite for models that give them some leverage on long-term problems or new ways of thinking about old problems.



Rob Axtell



A Pashtun wedding.


For example, the epidemiology community, driven mostly by demands that arose following 9/11, is using agent-based models to determine the best U.S. response to bioterrorism. Who gets vaccinated and in what order? Should you shut down interstate trucking or all schools and universities?

Agent-based traffic models have also gained a foothold. Models of the entire transportation system give you much higher fidelity than do equation-based simulations, and the policymakers want that because they need to know what to do.

In Washington there's also a lot of interest in modeling certain kinds of military operations, both logistical movements and person-to-person combat, in which every weapon the soldier will have is represented.

The models simulate various military rules of engagement to

In this simulation of Axtell's retirement model, there are 100 agents (squares) for each age from 19 to 100 (only ages 55 and above are shown). Purple agents are "rationals" and retire at age 65, the earliest possible retirement age. Blue agents are "imitators" who imitate the people in their social network. Yellow agents are "randoms," they have some probability of retiring at 65 or later. Red agents are retired, white agents have died. Initially (left), retirement at age 65 is not yet the "norm," and lots of agents retire after age 75. Over time (right), 65 becomes the norm as imitators follow rationals.



determine which are likely to produce successful scenarios, and the results are decision aids for policymakers.


1663: Do you think decision makers have a realistic view of what models can do?


Axtell: In some cases, they may be jumping the gun. They want substantive and quantitative results, but you can't get such results from agent modeling without first gathering data on the agents.

MacKerrow: Yes. Our Washington customers often ask us to do what I call general models. What causes terrorism, and will we see more? Will Al-Qaeda form into new Al-Qaedas? We've had little success in answering those questions because terrorism is used for different reasons in different contexts.

However, if we focus on specific cases, for example, a particular tribe of Pashtuns [an ethno-linguistic group with populations in Afghanistan] and have good data on what they've been doing for the last 50 years, we may be able to simulate what might happen to the tribe in the future.

For example, suppose a new warlord is suddenly running the village. We could address specific questions: How might that warlord affect opium farming in the region, and how might changes in opium production affect that tribe's relative regional power? If we couldn't talk to the tribal members directly, we would work with anthropologists who have studied that tribe or similar tribes. When we focus on specific situations, with knowledge of the relevant behavioral drivers, agent-based simulation makes sense.





The Pashtuns follow patrilineal marriage rules, with greater age conferring greater honor (age increases toward the right). Ideally, the oldest son (far right) would marry a daughter of his oldest uncle (blue arrows) but the pink arrow indicates his desire to marry the cousin who has no sisters and is likely to have a larger dowry. Marrying her, however, would dishonor the oldest uncle's family, possibly leading to violent feuds. Rivalry between male cousins causes many tribes to split into sub-tribes ("khels").

Axtell: Yes, making a good model is always hard. You have to know your sources and the people you're trying to model. I spent 5 years working on a model of 1,000 Anasazi Native Americans who lived in a northeastern Arizona valley from 500 A.D. to 1200 A.D. We had to learn a lot about their archaeology and anthropology.

1663: And what question were you answering?

Axtell: We were trying to adjudicate a longstanding debate about the demise of the Anasazi. Was it due primarily to microclimatic fluctuations or to social factors like war, an epidemic, or even a religious cult?

Our model showed that most of the fluctuations in their population could be understood with the environmental story. The tree rings gave us a record of how much rain actually fell, and so by linking that to their environment and lifestyle in a credible way—the calories they needed to survive, the typical number of offspring per person, life expectancy, and so on—we could get the ebb and flow of the population about right without adding in factors like war or pestilence. The final decline of their last few years, with the last 50 people, probably involved sociality, but their overall decline could be explained by drought caused by microclimatic shifts.

MacKerrow: The nice thing about agent-based models is that you can run the simulations many times to see the macro results, such as total population, time after time, and then compare the results of all the runs with each other and with historical fact.

Axtell: In essence, history represents only one run of the model. You hope that if you run the model 1,000 times, a typical run will agree with history. Of course there's always the possibility of extreme events, like the birth of an Adolph Hitler on steroids who totally annihilates a population.

I always think of the agents as abstract particles with personality.

MacKerrow: But fundamental social science still can't explain why someone born in London ends up as a suicide bomber in Iraq. We sort of understand how attitude diffusion occurs, and we're learning about suicide bombing, but we're still in the dark about the basic drivers behind the behavior. It's like living in a time before the apple fell on Newton's head and being asked to analyze thousands of bits of video data on the trajectories of baseballs. You need Newton's laws of mechanics before you can really explain those trajectories.

Axtell: Right. The social sciences are really pre-Newtonian. We don't even know the right units of analysis. If we actually knew how the 25 or 50 main components of the brain work and interact, we could build a model that would have deep behavioral significance. We don't have that information yet.

MacKerrow: All is not grim though. We are really in a renaissance period for the social sciences. Vast amounts of attitude data are sitting out on the Internet. We have empirical observations from experimental economics studies. Functional magnetic resonance imaging even allows us to directly observe a thinking brain in action. If all that information can be coupled to the experimental analysis of computational social science, we could see great advancements.

The major national security issues today revolve on worldwide social-cultural dynamics. Which group will become violent next? Which coalitions will form or break up? Which social identity groups will support, oppose, or react violently to proposed policies? Agent-based simulation may provide a new "algebra" to better represent the way people really behave. ❖

SPOTLIGHT


Better Breast Cancer Detection and Diagnosis

Women may eventually have access to safer, more-comfortable, and more-accurate breast cancer scans. Currently, the only routine breast-screening technology is mammography, which uses low-dose x-rays to scan through tissue and capture on film a two-dimensional (2D) projection of the breast.

Los Alamos scientist Lianjie Huang, in collaboration with researchers from Karmanos Cancer Institute (KCI), London's Imperial College, and Stanford University, has developed a better way, producing a three-dimensional (3D) image, using not x-rays, but sound waves.

The technique, called ultrasound computed tomography (ultrasound CT), uses a prototype scanning device built at KCI. A woman's breast is immersed in water and surrounded by a ring-shaped array of hundreds of ultrasound elements. Each element emits ultrasound waves and then receives waves that are scattered from the soft tissue. The array is moved incrementally down the entire breast, gathering data at each step.

A suite of newly developed computer algorithms converts the stepwise ultrasound



This view of a 3D ultrasound CT image was obtained, from a patient, using KCI's prototype device. It shows a cross section of the breast near the chest wall (top of image) and a vertical cross section through the remainder of the breast. A tumor (red) is visible near the chest wall.

data into a series of high-resolution, 2D images and then turns the series into a single 3D image. The technology actually obtains three kinds of images, corresponding to the speed, attenuation, and reflectivity of the waves. The wealth of information allows ultrasound CT to single out cancerous lesions more accurately than can today's mammography.

Ultrasound CT has the potential to detect cancer in its earliest stages. And since it is both safer (no ionizing radiation) and more comfortable (there is no need to compress tissue as there is in mammography), it should prove an attractive alternative for future breast cancer screening.

Getting inside a Fly's Head

Why in the world would anyone care about a fly's head?

A team of researchers at Los Alamos (Ilya Nemenman), Princeton, and Indiana University has looked into one and learned something new about neurons, the electrically charged cells that transmit information in the nervous system. They've disproved the current view of how neurons transmit data in the part of a fly's brain that processes visual information.

Information is passed between neurons in a series of voltage "spikes." Because the spikes all have the same shape, they can convey information only through their placement in time. Scientists have long debated whether information is transmitted in the lengths of the time intervals between spikes or in the average number of spikes in much larger windows of time. The latter model has generally prevailed and has been used in the development of artificial neural networks—electrical circuits or software designed to mimic how biological nerve systems process sensory input.

This team found instead that a fly's visual system conveys information by controlling the intervals between successive spikes,



A fly's visual system can help us design better artificial neural networks.

IMAGE COURTESY OF DENNIS KUNKEL MICROSCOPY, INC

to a precision of 200 millionths of a second. The finding sheds light not only on how these and possibly other types of neurons actually work but also suggests a reason why existing artificial neural networks designed to, for example, recognize faces do not work very well.

The scientists developed a new mathematical technique to analyze the timing of the spikes produced by a neuron sensitive to horizontal motion in a fly's field of view. The fly was immobilized on a small programmable turntable spinning fast enough to simulate the horizontal-rotation component of the insect's aerial acrobatics. To provide extra realism, team members placed the apparatus outdoors, where the visual background was far more natural and complex than it could be in a laboratory and where the fly's reactions to natural fluctuations in light—caused by passing clouds, for example—could also be recorded and analyzed.

The team believes the laboratory settings of previous studies were too monotonous and predictable to allow access to accurate data about how the fly's visual system responds in the natural environment.

The Sound of One Star Falling

In 1908, a meteor several hundred meters in diameter exploded when it entered Earth's atmosphere over Russia. The blast wave flattened forests and jiggled the pens of newly invented barographs far away in Britain.

On a vastly smaller scale, a meteor with a diameter less than 10 centimeters (about the size of a softball) can also produce


atmospheric disturbances at the ground. A small meteor, frictionally heated by the air rushing past it, produces its own “blast wave,” a pulse of low-frequency sound waves (infrasound). The frequency is 0.02 to 20 hertz—too low for humans to hear. In a collaborative project, the Laboratory’s Doug ReVelle and students and faculty at the University of Western Ontario are using infrasound to detect small meteors and estimate their energy, location, and duration.

The researchers are using an array of infrasonic sensors, as well as radar, video cameras, and seismic and radio sensors. The equipment is located in Ontario. Every month the scientists observe at least one meteor that can produce infrasound detectable at the ground—a flux at least 100 times higher than earlier observations had suggested. They have also found that meteor-infrasound theory—developed by ReVelle more than 30 years ago—agrees well with the measurements. The theory had been untestable until now because only a handful of infrasonic observations of small meteors had been made.

“Infrasound can also be used to observe large manmade chemical and nuclear explosions,” ReVelle says. “Although such explosions can be intentionally hidden from satellites, their ‘sound effects’ can still give them away.” In fact, the Department of Energy supports the infrasound monitoring program at Los Alamos, the only such program in the United States. The program operates six infrasonic-sensor arrays in western states. The arrays are routinely used to monitor White Sands Missile Range test explosions, NASA Space Shuttle launches and reentries, smaller missile launches, earthquakes, volcanic eruptions, and gas-fire explosions, among other events.



A sensor used to detect small meteors.



Sandrine Gaillard, a Laboratory affiliate, stares into Trident’s north target chamber where ultra-short laser pulses create high-energy-density plasmas as well as monoenergetic ion pulses for cancer research applications.

Exploits with Superbright Light

Amazing things can be done when the small amount of energy needed to light a 100-watt bulb for one second is packed into a tiny pulse of laser light lasting but a trillionth of a second and having a spot size about 100 times smaller than the period at the end of this sentence. That’s what Los Alamos scientists have demonstrated with the new ultra-short, highly-focused, ultra-intense pulses available from the newly enhanced Trident laser, now delivering power at one-tenth of a petawatt (a petawatt equals a million-billion watts).

Shine one of those tiny pulses on a thin foil target and it produces an intense pulse of x-rays with the right energies (18–35 thousand volts—kilovolts) to make high-precision images of imploding fusion capsules, just what will be needed to diagnose fusion experiments at the National Ignition Facility at Lawrence Livermore National Laboratory. Prepare that foil target in a slightly different way, and the Trident laser pulse will produce a beam of 50-million-electronvolt protons, a beam with 10 times more energy content than found in proton beams from similar lasers for the same laser intensity.

At this energy and intensity the proton

beam has a host of potential uses, from making images (proton radiographs) of dense objects to making radioisotopes for medical applications or performing tabletop nuclear physics experiments. Slightly higher energies make possible other applications, like treating cancers or screening containers for the presence of plutonium or uranium. The laser itself might even be used to divert lightning from power lines and buildings.

With all these applications in sight, Kirk Flippo, one of the Trident physicists, is excited about the future. “Short-pulse technology could revolutionize many aspects of our lives, and Trident is one of a handful of systems leading the way.”

Other scientists emphasize Trident’s basic research potential. It can re-create states of matter seen only around black holes and inside gamma-ray bursts, thereby expanding the field of laboratory astrophysics. It can also produce ultra-short x-ray pulses to image the gyrations of proteins as they fold and can even be used to observe such bizarre phenomena as “Unruh radiation,” the theoretical radiation emitted by a particle subjected to massive acceleration.

Trident has just become a national user facility, opening its doors to scientists everywhere and giving all a chance to channel their ideas into real experiments.

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by Los Alamos National Security, LLC, for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396.

A U.S. Department of Energy Laboratory LALP-08-026



Address mail to

1663

Mail Stop M711

Los Alamos National Laboratory

Los Alamos, NM 87545

email: 1663magazine@lanl.gov

Tel: 505-667-1447

Fax: 505-665-4408

www.lanl.gov/science/1663/



Los Alamos, the city on the hill, viewed across the valley from the Santa Fe ski area.

*PRINCIPAL ASSOCIATE DIRECTOR
SCIENCE, TECHNOLOGY, AND ENGINEERING—TERRY WALLACE*

EDITOR-IN-CHIEF AND WRITER—JAY SCHECKER

SCIENCE EDITOR AND WRITER—NECIA GRANT COOPER

MANAGING EDITOR AND WRITER—EILEEN PATTERSON

GRAPHIC DESIGNER—GLORIA SHARP

ILLUSTRATOR—DONALD MONTOYA

PHOTOGRAPHER—LEROY N. SANCHEZ

CONTRIBUTING WRITERS—BRIAN FISHBINE, CRAIG CARMER

DISTRIBUTION COORDINATORS—JEANETTE GALLEGOS, TAMMY MILLIGAN

PRINTING COORDINATOR—LUPE ARCHULETA

ADVISORY BOARD—JEFF BERGER, JAMES RICKMAN, KIM THOMAS